

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (Currently amended) A method for assigning traffic buckets to a cache system, the method comprising:

a) when a new cache system starts up in a cache cluster having a plurality of cache systems among which a plurality of total buckets are to be allocated, determining a full bucket allocation for the new cache system;

b) periodically determining a load of the new cache system;

c) when it is determined that the new cache system is underloaded, slowly assigning a portion of the full bucket allocation or a portion of previously shed buckets to the new cache system unless the full bucket allocation has already been assigned to the new cache system; and

d) when it is determined that the new cache system is overloaded ~~overloaded~~, shedding a portion of the buckets previously assigned to the new cache system.

2. (original) A method as recited in claim 1 further comprising assigning the full bucket allocation to the new cache system when the cache cluster is operating at a maximum load.

3. (previously presented) A method as recited in claim 1 wherein slowly assigning a portion of the full bucket allocation to the new cache comprises:

initially assigning a portion of the full bucket allocation to the new cache system;

when no buckets have been previously shed, assigning a portion of the unassigned buckets to the new cache system; and

when buckets have been previously shed, assigning a portion of a number of buckets that were previously shed from the new cache system.

4. (previously presented) A method as recited in claim 1, wherein shedding a portion of the buckets previously assigned to the new cache comprises:

when no buckets have been previously shed, periodically shedding a portion of the assigned buckets from the new cache system;

when buckets have been previously shed, periodically shedding a portion of a number of buckets that were previously shed from the new cache system;

5. (previously presented) A method as recited in claim 4 further comprising after the full allocation of buckets is assigned to the new cache system or the last assigned or shed number of buckets for the new cache system is less than or equal to one:

halting operations (b) through (d);

e) periodically monitoring a load of each of the other cache systems, including the new cache system, within the cache cluster;

f) when any of the cache systems is overloaded, shedding the minimum number of buckets from the overloaded ~~other~~ cache system; and

g) when any of the cache systems is underloaded, adding the minimum number of buckets to the underloaded ~~other~~ cache system.

6. (previously presented) A method as recited in claim 5, wherein the minimum number of buckets equals a single bucket, the portion of the unassigned buckets equals a half of the unassigned buckets, and the portion of a number of buckets that were previously shed equals half of the number of buckets that were previously shed.

7. (previously presented) A method as recited in ~~claim 4~~ claim 1 wherein the ~~second~~ technique operations (b) through (d) are is performed until the full allocation has been assigned

to the new cache system or a minimum number of buckets have been added to or shed from the new cache system.

8. (previously presented) A method as recited in claim 7 wherein operations (b) through (d) are performed about every 30 seconds.

9. (Cancelled)

10. (previously presented) A method as recited in claim 1 further comprising:
when an existing cache system leaves the cache cluster or shuts down,
determining a new bucket allocation for each of the remaining cache systems; and
assigning buckets to the remaining cache systems using operations (e) through (g).

11. (previously presented) A method as recited in claim 5 further comprising:
when an existing cache system leaves the cache cluster or shuts down,
determining a new bucket allocation for each of the remaining cache systems; and
assigning buckets to the remaining cache system using the new bucket allocation with operations (e) through (g).

12. (original) A method as recited in claim 1, wherein the full bucket allocation is equal to a number of buckets allocated to each existing cache system within the cache cluster.

13. (original) A method as recited in claim 1, wherein the full bucket allocation is not equal to a number of buckets allocated to each existing cache system within the cache cluster.

14. (original) A method as recited in claim 13, further comprising receiving a weight value from the new cache system indicating a percentage of the total buckets to allocate the new cache system.

15. (previously presented) A method as recited in claim 4 further comprising:

receiving load information from the new cache, the load information indicating whether the new cache system is overloaded; and

using the load information to determine whether the new cache is overloaded.

16. (original) A method as recited in claim 15 wherein load information is periodically received from the new cache system.

17. (original) A method as recited in claim 15 wherein the load information further indicates a number of buckets to shed or add, wherein the portion of the number of buckets that were previously shed from the new cache system, the portion of the unassigned buckets, and the portion of the assigned buckets are equal to the indicated number of buckets to shed or add.

18. (Currently amended) A computer system operable to assign traffic buckets to a cache system, comprising:

a memory; and

a processor coupled to the memory,

a) when a new cache system starts up in a cache cluster having a plurality of cache systems among which a plurality of total buckets are to be allocated, determining a full bucket allocation for the new cache system;

b) periodically determining a load of the new cache system;

c) when it is determined that the new cache system is underloaded, slowly assigning a portion of the full bucket allocation or a portion of previously shed buckets to the new cache system unless the full bucket allocation has already been assigned to the new cache system; and

d) when it is determined that the new cache system is overloaded
~~overloaded~~, shedding a portion of the buckets previously assigned to the
new cache system.

19. (original) A computer system as recited in claim 18, wherein at least one of the
memory and the processor are further adapted to provide:

assigning the full bucket allocation to the new cache system when the
cache cluster is operating at a maximum load.

20. (previously presented) A computer system as recited in claim 18, wherein at least
one of the memory and the processor are further adapted to provide after the full allocation of
buckets is assigned to the new cache system or the last assigned or shed number of buckets for
the new cache system is less than or equal to one:

halting operations (b) through (d);

e) periodically monitoring a load of each of the other cache systems,
including the new cache system, within the cache cluster;

f) when any of the cache systems is overloaded, shedding the minimum
number of buckets from the overloaded cache system; and

g) when any of the cache systems is underloaded, adding the minimum
number of buckets to the underloaded cache system.

21. (previously presented) A computer system as recited in claim 18, wherein the e
operations (b) through (d) are is performed until the full allocation has been assigned to the new
cache system or a minimum number of buckets have been added to or shed from the new cache
system and slowing assigning comprises (i) initially assigning a portion of the full bucket
allocation to the new cache system; (ii) when no buckets have been previously shed, periodically
assigning a portion of the unassigned buckets to the new cache system; and (iii) when buckets

have been previously shed, periodically assigning a portion of a number of buckets that were previously shed from the new cache system, and wherein shedding a portion of the buckets previously assigned to the new cache system comprises (i) when no buckets have been previously shed, periodically shedding a portion of the assigned buckets from the new cache system; and (ii) when buckets have been previously shed, periodically shedding a portion of a number of buckets that were previously shed from the new cache system;

22. (original) A computer system as recited in claim 21 wherein the portion of the number of buckets that were previously shed from the new cache system, the portion of the unassigned buckets, and the portion of the assigned buckets are equal to a half portion.

23. (currently amended) A computer system as recited in claim 21 wherein at least one of the memory and the processor are further adapted to provide:

when an existing cache system leaves the cache cluster or shuts down,
determining a new bucket allocation for each of the remaining cache systems; and
assigning buckets to the remaining cache system ~~the first technique~~
operations (e) through (g).

24. (previously presented) A computer system as recited in claim 18 wherein at least one of the memory and the processor are further adapted to provide:

when an existing cache system leaves the cache cluster or shuts down,
determining a new bucket allocation for each of the remaining cache systems; and
assigning buckets to the remaining cache system using the new bucket
allocation with operations (e) through (g).

25. (Currently amended) A computer program product for assigning traffic buckets to a cache system, the computer program product comprising:

at least one computer readable medium;

computer program instructions stored within the at least one computer readable product configured to:

a) when a new cache system starts up in a cache cluster having a plurality of cache systems among which a plurality of total buckets are to be allocated, determining a full bucket allocation for the new cache system;

b) periodically determining a load of the new cache system;

c) when it is determined that the new cache system is underloaded, slowly assigning a portion of the full bucket allocation or a portion of previously shed buckets to the new cache system unless the full bucket allocation has already been assigned to the new cache system; and

d) when it is determined that the new cache system is overloaded ~~overerloaded~~, shedding a portion of the buckets previously assigned to the new cache system.

26. (previously presented) A computer program product as recited in claim 25, wherein the computer program instructions are further configured to assign the full bucket allocation to the new cache system when the cache cluster is operating at a maximum load.

27. (previously presented) A computer program product as recited in claim 25, wherein slowing assigning a portion of the full bucket allocation to the new cache comprises:

initially assigning a portion of the full bucket allocation to the new cache system;

when no buckets have been previously shed, assigning a portion of the unassigned buckets to the new cache system; and

when buckets have been previously shed, assigning a portion of a number of buckets that were previously shed from the new cache system.

28. (previously presented) A computer program product as recited in claim 25, wherein shedding a portion of the buckets previously assigned to the new cache comprises:

when no buckets have been previously shed, periodically shedding a portion of the assigned buckets from the new cache system;

when buckets have been previously shed, periodically shedding a portion of a number of buckets that were previously shed from the new cache system.

29. (previously presented) A computer program product as recited in claim 28 wherein the computer program instructions are further configured to after the full allocation of buckets is assigned to the new cache system or the last assigned or shed number of buckets for the new cache system is less than or equal to one:

halt operations (b) through (d);

e) periodically monitor a load of each of the other cache systems, including the new cache system, within the cache cluster;

f) when any of the cache systems is overloaded, shed the minimum number of buckets from the cache system; and

g) when any of the cache systems is underloaded, add the minimum number of buckets to the underloaded cache system.

30. (previously presented) A computer program product as recited in claim 29, wherein the minimum number of buckets equals a single bucket, the portion of the unassigned buckets equals a half of the unassigned buckets, and the portion of a number of buckets that were previously shed equals half of the number of buckets that were previously shed.

31. (previously presented) A computer program product as recited in claim 25, wherein the operations (b) through (d) are is performed until the full allocation has been assigned

to the new cache system or a minimum number of buckets have been added to or shed from the new cache system

32. (previously presented) A computer program product as recited in claim 31, wherein operations (b) through (d) are performed about every 30 seconds.

33. (previously presented) A computer program product as recited in claim 25, wherein the computer program instructions are further configured to:

when an existing cache system leaves the cache cluster or shuts down,
determine a new bucket allocation for each of the remaining cache systems; and
assign buckets to the remaining cache systems using the new bucket allocation with operations (e) through (g).

34. (previously presented) A computer program product as recited in claim 29, wherein the computer program instructions are further configured to:

when an existing cache system leaves the cache cluster or shuts down,
determine a new bucket allocation for each of the remaining cache systems; and
assign buckets to the remaining cache system using the new bucket allocation with operations (e) through (g).

35. (previously presented) A computer program product as recited in claim 25, wherein the full bucket allocation is equal to a number of buckets allocated to each existing cache system within the cache cluster.

36. (previously presented) A computer program product as recited in claim 25, wherein the full bucket allocation is not equal to a number of buckets allocated to each existing cache system within the cache cluster.

37. (previously presented) A computer program product as recited in claim 36, wherein the computer program instructions are further configured to receive a weight value from the new cache system indicating a percentage of the total buckets to allocate the new cache system.

38. (previously presented) A computer program product as recited in claim 28, wherein the computer program instructions are further configured to:

receive load information from the new cache, the load information indicating whether the new cache system is overloaded; and

use the load information to determine whether the new cache is overloaded.

39. (previously presented) A computer program product as recited in claim 38, wherein load information is periodically received from the new cache system.

40. (previously presented) A computer program product as recited in claim 38, wherein the load information further indicates a number of buckets to shed or add, wherein the portion of the number of buckets that were previously shed from the new cache system, the portion of the unassigned buckets, and the portion of the assigned buckets are equal to the indicated number of buckets to shed or add.

41. (currently amended) An apparatus for assigning traffic buckets to a cache system, comprising:

means for when a new cache system starts up in a cache cluster having a plurality of cache systems among which a plurality of total buckets are to be allocated, determining a full bucket allocation for the new cache system;

means for periodically determining a load of the new cache system;

means for when it is determined that the new cache system is underloaded, slowly assigning a portion of the full bucket allocation or a portion of previously shed buckets to the new cache system unless the full bucket allocation has already been assigned to the new cache system; and

means for when it is determined that the new cache system is overerloaded, shedding a portion of the buckets previously assigned to the new cache system.